

CDB on pp. 13-15

"CDB" stands for "cache data base", which is apparent from page 13, line 14, which speaks of "cache data base description manager 303", later termed "CDB description manager 303".

DSID on p. 13

DSID stands for "data set identifier", as is apparent when page 14, line 13 is compared with page 14, line 11 and page 14, line 12. Applicants have amended page 13, line 34 through page 14, line 20 so that the reference numbers in the text track the reference numbers in FIG. 5.

Rcvr on page 14

"Rcvr" stands for "Receiver". Applicants' attorney believes that "rcvr" is a generally-used abbreviation for "receiver", and in any case, page 14, lines 21-26 describes "Update Rcvr 321" in sufficient detail so that there should be no doubt as to what it does.

The rejection of new claims 112-131 under 35 U.S.C. 112, first paragraph

Examiner rejects these claims because they "contain subject matter which was not described in the Specification". She bases her finding of lack of support on the following;

Claims 112-131 recite objects, distributed database system, redirector, or specifier/specifiers which are not mentioned in Applicant's Specification or Drawings.

When looked at in light of MPEP 2163-2163.05's discussion of the Written Description requirement of 35 U.S.C. 112, first paragraph, the rejection is clearly problematical. First, there is no requirement that the Specification contain the terminology used in the claims; instead,

... the fundamental factual inquiry is whether a claim defines an invention that is clearly conveyed to those skilled in the art at the time the application was filed. The subject matter of the claim need not be described literally (i.e., using the same terms or *in haec verba*) in order for the disclosure to satisfy the description requirement. If a claim is amended to include subject matter, limitations, or terminology not present in the application as filed, *involving a departure from, addition to, or deletion from the disclosure as filed*, (emphasis added), the examiner should conclude that the claimed subject matter is not described in the application. (MPEP 2163.02)

Second, when Applicants filed their new claims 112-131 in their response of 3/12/02, they demonstrated that new claims 112-131 are fully supported by the disclosure of the above application as follows:

New claims 112-131 are fully supported by the disclosure of the present application and consequently add no new matter. Taking claim 112 as an example, the "request" of the claim is embodied in the queries of the present application, the "plurality of database systems" is at least embodied in cache database 347 and source database 241, with cache database 347 embodying the "first database system" and source database 241 embodying the "second database system". The "redirector" is embodied in DA interface 304 and query dispatcher 351. The manner in which these components of this embodiment of the invention interact to "respond[]" to the request when the request includes a specifier that cannot be interpreted in the first database system by causing the request to be executed at least in part in a second database system of the plurality, the request otherwise being executed in the first database system" is described at least at page 12, lines 15-24. (Applicants' response of 3/12/02, p. 16)

According to MPEP 2163.04,

If applicant amends the claims and points out where and/or how the originally filed disclosure supports the amendments(s), and the examiner finds that the disclosure does not reasonably convey that the inventor had possession of the subject matter of the amendment at the time of the filing of the application, the examiner has the initial burden of presenting evidence or reasoning to explain why persons skilled in the art would not recognize in the disclosure a description of the invention defined by the claims. *Accordingly, the examiner should identify what portions of the amendment lack support in the originally filed disclosure and should fully explain the basis for the examiner's finding. The examiner should also comment on the substance of the applicant's remarks.* (emphasis added)

Since Applicants have already demonstrated that the Specification as originally filed supports their new claims 112-131, MPEP 2163.04 clearly places the burden on Examiner to "identify what portion(s) of the amendment lack support", and Applicants' attorney respectfully submits that Examiner's rejection of claims 112-131 under 35 U.S.C. 112, first paragraph has not met this burden and is therefore without basis.

The rejection of claims 24-35, 53-60, and 84-93 under 35 U.S.C. 103(a) as unpatentable over Draper

Prior rejection based on Draper

Claims 24 and 34 were rejected under 35 U.S.C. 102(e) as anticipated by this reference in the Office action mailed 6/20/00 in the above application. Applicants traversed the rejection in their response mailed 10/20/00 in the application.

5

Examiner's failure to make a prima facie case for the rejection of claims 24, 34, and 84

As set forth at MPEP 2142, a *prima facie* case for a rejection under 35 U.S.C. 103 must meet three criteria:

10 First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference . . . must teach or suggest all the claim limitations. (MPEP 2142)

15 As will be shown in the following, Draper does not disclose what Examiner believes it does, and she has consequently not made her *prima facie* case. Claim 24 is exemplary for what Applicants are claiming:

20 24. (twice amended) An improved server of the type that provides a program with a standard interface for querying remote datasets, the improved server comprising:
a queryable cache that contains copies of certain of the datasets and is local to the server,
the improved server receiving a query for a remote dataset in a form required by the interface from the program, determining whether a copy of a dataset to be queried is
25 present in the queryable cache, and, if the copy is present, querying the copy, and otherwise querying the remote dataset,
whereby the queryable cache is transparent to the program.

30 The claim is addressed to an improved *server of the type that provides a program with a standard interface for querying remote datasets*. The server comprises a *queryable cache that contains copies of certain of the datasets and is local to the server*. The server further does three things:

1. it receiv[es] a query for a remote data set in a form required by the interface from the program;
2. it determin[es] whether a copy of the data set to be queried is present in the queryable
35 cache; and

3. *if the copy is present, it quer[ies] the copy and otherwise quer[ies] the remote data set.*

Beginning with the preamble, what is meant by a "server of the type that provides a program with a standard interface for querying remote datasets" is explained at page 2, lines 16-24 of Applicants' Specification:

5 Data access layer 112 is generally provided by the manufacturer of database server 115. It takes queries written in standard forms such as OLE-DB, ODBC, or JOBC, converts the queries into the form required by database server 115, and places the queries in messages in the form required by network 113. Database server 115 then
10 executes the query and returns the result via network 113 to data access layer 112, which puts the results into the required standard form and returns them to Web application 111, which in turn puts the result into the proper format for HTML program 109. HTML program 109 then uses the result in making the HTML page 106 to be returned to browser 103.

15 there is no indication anywhere in Draper that Draper's servers 106 are servers "of the typethat provides a program with a standard interface for querying remote datasets", as required by the claim..

Continuing with the body of the claim, the server comprises "a queryable cache that contains copies
20 of certain of the datasets and is local to the server". The first problem here is that as stated at col. 8, lines 11 and 12 of Draper, "The system 600 also includes two client caches 608, 610, which reside on clients 110". Draper's caches are thus not "local to the server", as required by the claim.

The next problem is that it is does not appear that Draper's caches are "queryable" and contain
25 "copies of certain of the datasets". What is meant by "queryable" and "datasets" is set forth at page 9, lines 11-19 of Applicants' Specification:

30 Because cached data 223 is contained in cache database 236, cached data 223 is *queryable*, that is, if a dataset is contained in cached data 223, queryable cache 219 can return as a result not only the entire dataset, but any subset of that dataset that can be described by a query. For example, if cached data 223 includes a dataset that lists all of the kinds of shirts sold by a company engaged in Web commerce and the list of kinds includes the colors that each kind of shirt is available in, queryable cache 219 will be able to handle a query for which the result is a list of the kinds of

shirt that are available in red.

5 A "queryable cache" is thus a cache in which queries can be run on the datasets contained in the cache. It is therefore neither a cache in which *results* of previous queries are stored and then returned when the identical query is repeated, nor is it a cache that stores data items such as Web pages that are simply returned in their entireties. There is nothing in Draper from which it can be concluded that Draper's cache is queryable. Examiner cites Draper, col. 2, lines 47-55 and col. 4, lines 23-28 in support of the proposition that Draper's cache is queryable, but col. 2, lines 57-55 describes indexes into tags, which are mechanisms for updating database systems and have nothing
10 to do with queries, and a distributed database, which is of course not a queryable cache. Col. 4, lines 23-28 merely disclose that Draper's system has network servers and clients; again, there is nothing about queryable caches.

15 There is also some indication in Draper that the caches are not queryable. There is first their location in client PC's and workstations; what makes a queryable database cache in a Web server desirable is that queries from many different users may be run on the cached datasets; when the cache is in a system where the queries are made by a single user, caching the datasets and running queries on them is far less efficient than simply caching the query results. This analysis is confirmed by the fact that what is cached in the exemplary embodiment described at col. 9, lines 33-45 is documents that
20 have been converted to HTML formats, i.e., objects that cannot be queried. Finally, in FIG. 6, the caches contain "cached data" and the discussion of FIG. 6 at col. 8, lines 1-21 speaks of "data items" in a fashion which seems to indicate that they are the contents of individual fields in data sets, and thus not queryable, rather than queryable objects such as datasets.

25 The remainder of the claim is directed to the interaction between the server and the queryable cache. Concerning the first part of the interaction, "the improved server receiving a query for a remote dataset in a form required by the interface from the program", there is simply no disclosure whatever in Draper about the interfaces used in the servers to query the database. Examiner cites col. 1, lines

53-55, but that merely speaks of how the system decides whether to use the cache or the disk in response to a database query or update operation. Col. 4, lines 23-28 and 61-67 and col. 5, lines 1-2 merely provide a general description of the servers and network clients making up Draper's system. Examiner herself concedes that the remaining elements of the interaction are not disclosed in Draper.

5

Since Draper does not in fact disclose the limitations of Applicants' claim 24 which Examiner takes the reference to disclose, Draper cannot serve as the basis for Examiner's rejection of claim 24; as Examiner will recognize, the same logic applies to the other independent claims 34 and 84.

10 *Rejections of the dependent claims*

As just pointed out, independent claims 24, 24, and 84 are all patentable over Draper; that being the case, all of the claims dependent from these claims are also patentable. It should, however, also be pointed out that the portions of Draper which Examiner relies on for rejecting the dependent claims also do not disclose the added limitations of the dependent claims, and that the dependent claims are

15 consequently patentable in their own rights.

Many of the rejections of the dependent claims are based on the portions of Draper which disclose the indexed tags which Draper uses to keep track of the types and order of operations on the data items in the database system. Tags are defined as follows at col. 5, lines 19-34:

20 Each data item 202 has an associated tag 204 . . . Each tag 204 value corresponds to an event in the history of the associated data item 202, such as the most recent update to the data item 202. The tags 204 may be designed to allow recovery of earlier versions of the data item 202. The tags 204 may also allow the system to determine which copy of a data item 202 is more recent,

25 so that the most recent data can be propagated during synchronization. . . . Suitable tags 204 include timestamps, version numbers, sequence numbers, update reference numbers, transaction counters, and other means of determining the relative order of operations on data items 202.

30 As further pointed out at col. 5, lines 41-42, the difference between the tags of Draper and prior-art tags is that Draper's tags are indexed.

In her rejections of claims 25, 35, and 85, Examiner equates Draper's tags to Applicants' global and local identifiers. The latter are defined as follows at page 11, lines 5-14 of Applicants' Specification:

5 In preferred embodiment 301, Web application 111 uses global data set identifiers in queries. The Web applications 111 in all of the servers 203 use the same set of global data set identifiers. A cache data base 347 in a given server 203 has its own set of local data set identifiers for the data sets cached in cache data base 347. In preferred embodiment 301, then, one may speak of
10 *global* queries and query contexts that use global data set identifiers and *local* queries and query contexts that use local data set identifiers. In the preferred embodiment, query analyzer 313 uses cached data base description 305 to translate global query contexts into local query contexts.

15 As is apparent from the foregoing, Applicants' global and local identifiers are used in queries and have nothing whatever to do with Draper's tags. Consequently, Draper's disclosure of tags cannot render claims 25, 35, and 85 obvious.

Continuing with claims 26, 53, and 93, the further limitation in those claims further specifies that
20 "the query analyzer further indicates to the server whether the copy of the dataset is in the queryable cache". Examiner cites Draper, col. 7, lines 9-20 as showing this limitation. The cited location is, however, a discussion of event lists. An event list is defined as follows at col. 7, lines 10-13:

Each element in an event list contains enough information to allow a cache manager to recreate the corresponding event, in order to duplicate the effect
25 of the event and thus update the cache.

Clearly, none of this has anything to do with determining whether a particular copy of the dataset is in the queryable cache, and thus, the additional limitation of claims 26, 53, and 93 is also not disclosed in Draper.

30 With regard to claims 29, 30, 31, 32, 56, 58, 90, and 91, all of these claims set forth various aspects of a data set manager that

determines whether to add or remove a dataset by determining a likelihood that a query will be made to the dataset. (claim 29)

5 There is simply nothing in Draper about a "likelihood that a query will be made" or about using such a likelihood to determine whether to add or remove a dataset. The location cited by Examiner in her rejection, col. 6, lines 15-21, simply discloses how tags are associated with data items. Claim 30 further recites how the data manager may use a query log to determine the likelihood. In rejecting this claim, Examiner cites col. 6, lines 15-21 again and the description of Draper's update log at col. 10, lines 22-53. The update log is, however, not a log of queries, but a log of information that is
10 needed to synchronize data items in a data base with data items in a cache or a replica of the database.

Concerning claims 31, 32, 58, 90, and 91, Examiner confuses "event" as used in Draper with "event" as used by Applicants. In Draper, "each tag value corresponds to an event in the history of the
15 associated data item 202, such as the most recent update of the data item 202." (col. 5, line 22). What is meant by "event" in Applicants' claims is apparent from the following passage from Applicants' Specification:

20 For example, if a company engaged in Web commerce is having a 1-day sale on certain items for which there are datasets in source database 241, query information 208 may indicate the datasets for the items and the time of the 1-day sale. Using that information, DSM 213 can obtain the datasets from source database 241 and cache them in cache database 236 before the beginning of the sale and remove them from cache database 236 after the end of the sale. (p. 10, lines 5-10)

25 Clearly, "event" as used in Applicants' claims has only the remotest connection with "event" as used by Draper. Examiner herself admits that Draper does not teach the limitations of claims 33, 60, and 92.

Basing a rejection under 35 U.S.C. 103 on a single reference

30 In the above rejections, Examiner employs a single reference in a rejection under 35 U.S.C. 103 and supplies the limitations that are not disclosed in Draper by simply stating that inclusion of the missing

limitations would have been obvious. Thus, at section 7 of her Office action, Examiner states:

it would have been obvious to one having ordinary skill in the art at the time the invention was made to determine whether a copy of the data set is present in the queryable cache and if the copy is present, querying the copy, and otherwise querying the remote data set.

This mode of rejection telescopes the requirement that the reference(s) show all of the claimed limitations with the requirement that either the reference or the prior art suggest combining the limitations as they are combined in the claim. In particular, it is not clear what Examiner is relying on for the limitations that are not present in Draper. Should she in fact be relying on official notice of scientific theory or common knowledge in the art, Applicants' attorney respectfully requests that Examiner cite a reference that illustrates the theory or shows the common knowledge in the art, as set forth in MPEP 2144.03. Applicants' attorney further respectfully calls Examiner's attention to the statement in MPEP 2144.03, citing *In re Ahlert*, USPQ 418, 420-421 (CCPA 1970) that

"the facts so noticed serve to 'fill the gaps' which might exist in the evidentiary showing" and should not comprise the principle evidence upon which a rejection is based.


Conclusion

Applicants have been fully responsive to Examiner's non-final Office action mailed 6/4/02. Applicants have replaced the present informal Drawing with a formal Drawing. Applicants have further amended and explained their Specification as required to overcome Examiner's objection thereto. No new matter has been added to the Specification by way of the amendments to the Drawing or Specification. Applicants have further traversed Examiner's rejection of claims 112-131 under 35 U.S.C. 112, first paragraph, and have traversed Examiner's rejection of claims 24-35, 53-60, and 84-93 under 35 U.S.C. 103(a) as unpatentable over Draper.

Since Applicants have been fully responsive as provided in 37 C.F.R. 1.111(b), Applicants respectfully request that Examiner reconsider her rejections of the claims and pass the application to issue, as provided in 37 C.F.R. 1.111(a)(1). No fees are believed to be required by way of this

election. Should any be required, please charge them to deposit account number 50315.

Respectfully submitted,



Attorney of record,

Gordon E. Nelson

57 Central St., P.O. Box 782

Rowley, MA, 01969,

Registration number 30,093

Voice: 1-978-948-7632

Fax: 1-866-723-0359

9/3/02

Date


**Specification amendment with changes marked**

FIG. 5 shows details of CDB description 305. In a preferred embodiment, it is a table which has at least an entry 501 for each dataset of source database 241 of which here is a copy in cache database 347. Each entry 501 contains the global dataset identifier 503 for the data set, by which the dataset is known in all servers 107 with queryable caches 219 containing copies
5 of the dataset, the local data set identifier 505, by which the dataset is known in cache database 347, and number of queries 507, which indicates the number of times the dataset has been queried over an interval of time. In the preferred embodiment, number of queries 507 embodies query information 307.

- 10 An entry 501(i) for a given dataset is accessed in a preferred embodiment by a hash function 503, which takes global dataset ID ~~507~~511 for the dataset and hashes it into an entry index 509 in table 305. CDB description manager 303 then searches table 305 for the entry 501 whose field 503 specifies global DSID 511 beginning at entry index 509. If no such entry is found, the dataset is not in cache database 347 and CDB description manager 303 signals a
15 miss 311 to query analyzer 313. Table 305 may also include entries 501 for global datasets that are not presently cached in cache database 347; in such entries, local dataset ID 505 has a null value and a miss is returned in response to the null value. The purpose of such entries is to maintain number of queries information 507 for such data sets, so that dataset manager 323 can determine whether to add the entry's dataset to cache database 347.